

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	153	(description or summary or statement or excerpt) with (best adj match\$8)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:44
L2	66	1 and search\$7 and query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:49
L3	17	(description or summary or statement or excerpt) with (best adj match\$8) with (database or listing or table or dataset or report)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:59
L4	153	(description or summary or statement or excerpt) with (best adj match\$8)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:47
L5	1	4 and ((lookup or look-up) adj table) same (map or mapp\$6)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:48
L6	11	4 and ((lookup or look-up) adj table)and (map or mapp\$6)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:48
L7	11	4 and ((lookup or look-up) adj table) and (map or mapp\$6)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:48
L8	41	4 and search\$7 same query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:04
L9	15	8 and (database with (resid\$6 or stor\$6) with server)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:53
L10	7	8 and (database adj9 (resid\$6 or stor\$6) adj9 server)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:54

L11	80	(description or summary or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:01
L12	27	11 and search\$7 same query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:57
L13	17	(description or summary or statement or excerpt) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:00
L14	104	(description or summary or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:04
L15	1	14 and domain adj knowledge	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:03
L16	1	14 and (heuristic or domain) adj knowledge	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:03
L17	223	((heuristic or domain) near6 knowledge) with database	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:04
L18	1	17 and (description or summary or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:09
L19	1	18 and search\$7 same query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:05

L20	67	17 and search\$7 same query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:05
L21	1	20 and (description or summary or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:10
L22	104	(description or summary or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:14
L23	20	22 and client and server	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:11
L24	14	23 and web and internet	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:11
L25	107	(description or descriptive or summar\$9 or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:15
L26	16	25 and (rules! or heuristics!) and (map or mapp\$6)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:15

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	153	(description or summary or statement or excerpt) with (best adj match\$8)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:44
L2	66	1 and search\$7 and query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:49
L3	17	(description or summary or statement or excerpt) with (best adj match\$8) with (database or listing or table or dataset or report)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:59
L4	153	(description or summary or statement or excerpt) with (best adj match\$8)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:47
L5	1	4 and ((lookup or look-up) adj table) same (map or mapp\$6)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:48
L6	11	4 and ((lookup or look-up) adj table)and (map or mapp\$6)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:48
L7	11	4 and ((lookup or look-up) adj table) and (map or mapp\$6)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:48
L8	41	4 and search\$7 same query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:04
L9	15	8 and (database with (resid\$6 or stor\$6) with server)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:53
L10	7	8 and (database adj9 (resid\$6 or stor\$6) adj9 server)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:38

L11	80	(description or summary or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:01
L12	27	11 and search\$7 same query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 12:57
L13	17	(description or summary or statement or excerpt) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:00
L14	104	(description or summary or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:04
L15	1	14 and domain adj knowledge	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:03
L16	1	14 and (heuristic or domain) adj knowledge	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:21
L17	223	((heuristic or domain) near6 knowledge) with database	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:04
L18	1	17 and (description or summary or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:09
L19	1	18 and search\$7 same query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:05

L20	67	17 and search\$7 same query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:29
L21	1	20 and (description or summary or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:10
L22	104	(description or summary or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:14
L23	20	22 and client and server	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:29
L24	14	23 and web and internet	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:11
L25	107	(description or descriptive or summar\$9 or statement or excerpt or result or report or outlin\$5) with (best adj match\$8) with (database or listing or table or dataset or report or hyperlink\$6 or hyper-link\$6 or (text adj string))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:15
L26	16	25 and (rules! or heuristics!) and (map or mapp\$6)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:33
L27	79	((heuristic or domain) adj knowledge) same match\$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:21
L28	14	27 and best adj match\$6	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:27


L29	22094	search\$7 same query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:23
L30	34	29 and restaurant same seafood	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:28
L31	122	29 and restaurant same s(legal adj seafood)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:25
L32	1	29 and restaurant same (legal adj seafood)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:26
L33	1	restaurant same (legal adj seafood)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:26
L34	2	30 and best adj match\$6	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:28
L35	202	restaurant same seafood	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:28
L36	2	35 and best adj match\$6	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:28
L37	10416	best adj match\$6	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:28
L38	823	37 and search\$7 same query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:29

L39	357	38 and client and server	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:30
L40	220	39 and database with server	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:31
L41	95	40 and ((look\$3 adj up) or (lookup or look-up))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:31
L42	9	40 and ((look\$3 adj up) or (lookup or look-up)) adj table	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:32
L43	9	42 and match\$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:32
L44	5	43 and (rules! or heuristics!)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:33
L45	3	44 and (mapp\$5 or map)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:36
L46	28291	query near9 database	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:36
L47	325	46 and match\$5 near2 (rules! or heuristics!)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:37
L48	32	47 and (generat\$6 or creat\$6 or construct\$6 or build\$4) near9 query near9 rules!	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:38

L49	7	48 and (database adj9 (resid\$6 or stor\$6) adj9 server)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:40
L50	7	48 and (database adj9 ((resid\$6 or stor\$6) adj9 server))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:41
L51	7	48 and (database adj9 ((resid\$6 or store or stored or storing) adj9 server))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:45
L52	115	(search\$6 adj6 database) adj9 knowledge	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/05/09 13:45


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

[rules](#) [heuristic](#) [lookup table](#) [mapping](#) [search](#) [query](#) [database](#) [server](#)

Found 309 of 154,226

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)

Display results


[Search Tips](#)
[Try this search in The ACM Guide](#)
☐ Open results in a new window

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

 Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Query evaluation techniques for large databases](#)

Goetz Graefe

 June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

 Full text available: [pdf\(9.37 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

2 [MIL primitives for querying a fragmented world](#)

Peter A. Boncz, Martin L. Kersten

 October 1999 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 8 Issue 2

 Full text available: [pdf\(261.36 KB\)](#)

 Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)


In query-intensive database application areas, like decision support and data mining, systems that use vertical fragmentation have a significant performance advantage. In order to support relational or object oriented applications on top of such a fragmented data model, a flexible yet powerful intermediate language is needed. This problem has been successfully tackled in Monet, a modern extensible database kernel developed by our group. We focus on the design choices made in the Monet interpreter ...

Keywords: Database systems, Main-memory techniques, Query languages, Query optimization, Vertical fragmentation

3

[Special issue: AI in engineering](#)

D. Sriram, R. Joobhani
January 1985 **ACM SIGART Bulletin**, Issue 91

Full text available:  [pdf\(8.79 MB\)](#) Additional Information: [full citation](#), [abstract](#)

The papers in this special issue were compiled from responses to the announcement in the July 1984 issue of the SIGART newsletter and notices posted over the ARPAnet. The interest being shown in this area is reflected in the sixty papers received from over six countries. About half the papers were received over the computer network.

4 ObjectGlobe: Ubiquitous query processing on the Internet

R. Braumandl, M. Keidl, A. Kemper, D. Kossmann, A. Kreutz, S. Seltzsam, K. Stocker
August 2001 **The VLDB Journal — The International Journal on Very Large Data Bases**,
Volume 10 Issue 1


Full text available:  [pdf\(251.44 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

We present the design of ObjectGlobe, a distributed and open query processor for Internet data sources. Today, data is published on the Internet via Web servers which have, if at all, very localized query processing capabilities. The goal of the ObjectGlobe project is to establish an open marketplace in which *data* and *query processing capabilities* can be distributed and used by any kind of Internet application. Furthermore, ObjectGlobe integrates *cycle providers* (i.e., machi ...

Keywords: Cycle-, function- and data provider, Distributed query processing, Open systems, Privacy, Quality of service, Query optimization, Security

5 Optimal signature extraction and information loss


Christos Faloutsos, Stavros Christodoulakis
September 1987 **ACM Transactions on Database Systems (TODS)**, Volume 12 Issue 3

Full text available:  [pdf\(1.98 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Signature files seem to be a promising access method for text and attributes. According to this method, the documents (or records) are stored sequentially in one file ("text file"), while abstractions of the documents ("signatures") are stored sequentially in another file ("signature file"). In order to resolve a query, the signature file is scanned first, and many nonqualifying documents are immediately rejected. We develop a framework that includes primary key hashing, multiattribute hash ...

6 Searching in high-dimensional spaces: Index structures for improving the performance of multimedia databases

Christian Böhm, Stefan Berchtold, Daniel A. Keim
September 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 3

Full text available:  [pdf\(1.39 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

During the last decade, multimedia databases have become increasingly important in many application areas such as medicine, CAD, geography, and molecular biology. An important research issue in the field of multimedia databases is the content-based retrieval of similar multimedia objects such as images, text, and videos. However, in contrast to searching data in a relational database, a content-based retrieval requires the search of similar objects as a basic functionality of the database system ...

Keywords: Index structures, indexing high-dimensional data, multimedia databases, similarity search

7 Research sessions: query processing II: Efficient k-NN search on vertically decomposed data

Arjen P. de Vries, Nikos Mamoulis, Niels Nes, Martin Kersten

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available:  pdf(1.26 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Applications like multimedia retrieval require efficient support for similarity search on large data collections. Yet, nearest neighbor search is a difficult problem in high dimensional spaces, rendering efficient applications hard to realize: index structures degrade rapidly with increasing dimensionality, while sequential search is not an attractive solution for repositories with millions of objects. This paper approaches the problem from a different angle. A solution is sought in an unconvent ...

8 An analysis of XML database solutions for the management of MPEG-7 media descriptions

Utz Westermann, Wolfgang Klas

December 2003 **ACM Computing Surveys (CSUR)**, Volume 35 Issue 4

Full text available:  pdf(448.76 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

MPEG-7 constitutes a promising standard for the description of multimedia content. It can be expected that a lot of applications based on MPEG-7 media descriptions will be set up in the near future. Therefore, means for the adequate management of large amounts of MPEG-7-compliant media descriptions are certainly desirable. Essentially, MPEG-7 media descriptions are XML documents following media description schemes defined with a variant of XML Schema. Thus, it is reasonable to investigate curren ...

Keywords: MPEG-7, XML database systems, multimedia databases

9 Extracting predicates from mining models for efficient query evaluation

Surajit Chaudhuri, Vivek Narasayya, Sunita Sarawagi

September 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 3

Full text available:  pdf(698.37 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Modern relational database systems are beginning to support ad hoc queries on mining models. In this article, we explore novel techniques for optimizing queries that contain predicates on the results of application of mining models to relational data. For such queries, we use the internal structure of the mining model to automatically derive traditional database predicates. We present algorithms for deriving such predicates for a large class of popular discrete mining models: decision trees, nai ...

Keywords: Complex predicate optimization, simpler rules from complex predictive functions

10 External memory algorithms and data structures: dealing with massive data

Jeffrey Scott Vitter

June 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 2

Full text available:  pdf(828.46 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#)

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance

bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

11 SilkRoute: A framework for publishing relational data in XML

Mary Fernández, Yana Kadiyska, Dan Suciu, Atsuyuki Morishima, Wang-Chiew Tan
December 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 4

Full text available:  pdf(687.91 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

XML is the "lingua franca" for data exchange between interenterprise applications. In this work, we describe SilkRoute, a framework for publishing relational data in XML. In SilkRoute, relational data is published in three steps: the relational tables are presented to the database administrator in a canonical XML view; the database administrator defines in the XQuery query language a public, virtual XML view over the canonical XML view; and an application formulates an XQuery query over the publ ...

Keywords: XML, XML storage systems, XQuery

12 Object-based and image-based object representations

Hanan Samet
June 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 2

Full text available:  pdf(1.05 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

An overview is presented of object-based and image-based representations of objects by their interiors. The representations are distinguished by the manner in which they can be used to answer two fundamental queries in database applications: (1) Feature query: given an object, determine its constituent cells (i.e., their locations in space). (2) Location query: given a cell (i.e., a location in space), determine the identity of the object (or objects) of which it is a member as well as the re ...

Keywords: Access methods, R-trees, feature query, geographic information systems (GIS), image space, location query, object space, octrees, pyramids, quadrees, space-filling curves, spatial databases

13 The BEA streaming XQuery processor

Daniela Florescu, Chris Hillery, Donald Kossmann, Paul Lucas, Fabio Riccardi, Till Westmann, J. Carey, Arvind Sundararajan
September 2004 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 13 Issue 3


Full text available:  pdf(328.94 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

This paper describes the design, implementation, and performance characteristics of a commercial XQuery processing engine, the BEA streaming XQuery processor. This XQuery engine was designed to provide high performance for message-processing applications, i.e., for transforming XML data streams. The engine is a central component of the 8.1 release of BEA's WebLogic Integration (WLI) product. The BEA XQuery engine is fully compliant with the August 2002 draft of the W3C XML Query Language ...

Applying an information gathering architecture to Netfind: a white pages tool for a changing and growing Internet

Michael F. Schwartz, Calton Pu

October 1994 **IEEE/ACM Transactions on Networking (TON)**, Volume 2 Issue 5

Full text available:  [pdf\(1.71 MB\)](#)


Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)



15 Distributed operating systems

Andrew S. Tanenbaum, Robbert Van Renesse

December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

Full text available:  [pdf\(5.49 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Distributed operating systems have many aspects in common with centralized ones, but they also differ in certain ways. This paper is intended as an introduction to distributed operating systems, and especially to current university research about them. After a discussion of what constitutes a distributed operating system and how it is distinguished from a computer network, various key design issues are discussed. Then several examples of current research projects are examined in some detail ...



16 Anchor text mining for translation of Web queries: A transitive translation approach

Wen-Hsiang Lu, Lee-Feng Chien, Hsi-Jian Lee

April 2004 **ACM Transactions on Information Systems (TOIS)**, Volume 22 Issue 2

Full text available:  [pdf\(280.55 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

To discover translation knowledge in diverse data resources on the Web, this article proposes an effective approach to finding translation equivalents of query terms and constructing multilingual lexicons through the mining of Web anchor texts and link structures. Although Web anchor texts are wide-scoped hypertext resources, not every particular pair of languages contains sufficient anchor texts for effective extraction of translations for Web queries. For more generalized applications, the app ...


Keywords: Multilingual translation, anchor text mining, competitive linking algorithm, cross-language Web search, cross-language information retrieval



17 NSF workshop on industrial/academic cooperation in database systems

Mike Carey, Len Seligman

March 1999 **ACM SIGMOD Record**, Volume 28 Issue 1

Full text available:  [pdf\(1.96 MB\)](#)


Additional Information: [full citation](#), [index terms](#)



18 Scalable feature selection, classification and signature generation for organizing large text databases into hierarchical topic taxonomies

Soumen Chakrabarti, Byron Dom, Rakesh Agrawal, Prabhakar Raghavan

August 1998 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 7 Issue 3

Full text available:  [pdf\(281.37 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

We explore how to organize large text databases hierarchically by topic to aid better searching, browsing and filtering. Many corpora, such as internet directories, digital libraries, and patent databases are manually organized into topic hierarchies, also called *taxonomies*. Similar to indices for relational data, taxonomies make search and access more




efficient. However, the exponential growth in the volume of on-line textual information makes it nearly impossible to maintain such taxono ...

19 Index structures for selective dissemination of information under the Boolean model

Tak W. Yan, Héctor García-Molina

June 1994 **ACM Transactions on Database Systems (TODS)**, Volume 19 Issue 2

Full text available:  pdf(2.03 MB)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The number, size, and user population of bibliographic and full-text document databases are rapidly growing. With a high document arrival rate, it becomes essential for users of such databases to have access to the very latest documents; yet the high document arrival rate also makes it difficult for users to keep themselves updated. It is desirable to allow users to submit profiles, i.e., queries that are constantly evaluated, so that they will be automatically informed of new additions tha ...

20 On randomization in sequential and distributed algorithms

Rajiv Gupta, Scott A. Smolka, Shaji Bhaskar

March 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 1

Full text available:  pdf(8.01 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Probabilistic, or randomized, algorithms are fast becoming as commonplace as conventional deterministic algorithms. This survey presents five techniques that have been widely used in the design of randomized algorithms. These techniques are illustrated using 12 randomized algorithms—both sequential and distributed— that span a wide range of applications, including: primality testing (a classical problem in number theory), interactive probabilistic proof s ...

Keywords: Byzantine agreement, CSP, analysis of algorithms, computational complexity, dining philosophers problem, distributed algorithms, graph isomorphism, hashing, interactive probabilistic proof systems, leader election, message routing, nearest-neighbors problem, perfect hashing, primality testing, probabilistic techniques, randomized or probabilistic algorithms, randomized quicksort, sequential algorithms, transitive tournaments, universal hashing

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

 SEARCH

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Query evaluation techniques for large databases

Full text Pdf (9.37 MB)

Source [ACM Computing Surveys \(CSUR\)](#) [archive](#)
 Volume 25, Issue 2 (June 1993) [table of contents](#)
 Pages: 73 - 169
 Year of Publication: 1993
 ISSN:0360-0300

Author [Goetz Graefe](#) Portland State Univ., Portland, OR

Publisher ACM Press New York, NY, USA

Additional Information: [abstract](#) [references](#) [citations](#) [index terms](#) [review](#) [collaborative colleagues](#) [peer to peer](#)

Tools and Actions: [Discussions](#) [Find similar Articles](#) [Review this Article](#)
[Save this Article to a Binder](#) [Display Formats: BibTex](#) [EndNote](#)

DOI Bookmark: Use this link to bookmark this Article: <http://doi.acm.org/10.1145/152610.152611>
[What is a DOI?](#)

↑ ABSTRACT

↑ REFERENCES

Note: OCR errors may be found in this Reference List extracted from the full text article. ACM has opted to expose the complete List rather than only correct and linked references.

- 1 [Nabil R. Adam, John C. Worthmann, Security-control methods for statistical databases: a comparative study, ACM Computing Surveys \(CSUR\), v.21 n.4, p.515-556, Dec. 1989](#)
- 2 [I. Ahn, Richard Thomas Snodgrass, Partitioned storage for temporal databases, Information Systems, v.13 n.4, p.369-391, Oct. 1, 1988](#)
- 3 [Joseph Albert, Algebraic Properties of Bag Data Types, Proceedings of the 17th International Conference on Very Large Data Bases, p.211-219, September 03-06, 1991](#)
- 4 [Anastasia Analyti, Sakti Pramanik, Fast search in main memory databases, Proceedings of the 1992 ACM SIGMOD international conference on Management of data, p.215-224, June 02-05, 1992, San Diego, California, United States](#)
- 5 [David P. Anderson, Shin Tzou, G. S Graham, The DASH Virtual Memory System, University of California at Berkeley, Berkeley, CA, 1988](#)
- 6 [Gennady Antoshenkov, Dynamic Query Optimization in Rdb/VMS, Proceedings of the Ninth International Conference on Data Engineering, p.538-547, April 19-23, 1993](#)



[Subscribe](#) (Full Service) [Register](#) (Limited Service, Free) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

rules +heuristic +lookup table +mapping +search + query +d

SEARCH



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

rules heuristic lookup table mapping search query database server excerpt

Found 8 of 154,226

Sort results by

relevance



[Save results to a Binder](#)

[Try an Advanced Search](#)

Display results

expanded form



[Search Tips](#)

[Try this search in The ACM Guide](#)

☐ Open results in a new window

Results 1 - 8 of 8

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Exploiting early sorting and early partitioning for decision support query processing](#)

J. Claussen, A. Kemper, D. Kossmann, C. Wiesner

December 2000 **The VLDB Journal — The International Journal on Very Large Data**

Bases, Volume 9 Issue 3

Full text available: [pdf\(478.23 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

Decision support queries typically involve several joins, a grouping with aggregation, and/or sorting of the result tuples. We propose two new classes of query evaluation algorithms that can be used to speed up the execution of such queries. The algorithms are based on (1) *early sorting* and (2) *early partitioning*— or a combination of both. The idea is to push the sorting and/or the partitioning to the leaves, i.e., the base relations, of the query evaluation plans (QEPs) and ...

Keywords: Decision Support Systems, Early sorting and partitioning, Hash joins and hash teams, Performance evaluation, Query processing and optimization

2 [A software engineering perspective on algorithmics](#)

Karsten Weihe

March 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 1

Full text available: [pdf\(1.62 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

An algorithm component is an implementation of an algorithm which is not intended to be a stand-alone module, but to perform a specific task within a large software package or even within several distinct software packages. Therefore, the design of algorithm components must also incorporate software-engineering aspects. A key design goal is adaptability. This goal is important for maintenance throughout a project, prototypical development, and reuse in new, unforeseen context ...

Keywords: algorithm engineering

3 [Service discovery: Behavioral models as service descriptions](#)

Robert J. Hall, Andrea Zisman

November 2004 **Proceedings of the 2nd international conference on Service oriented computing**

Full text available:  pdf(209.70 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Interface descriptions, while adequate for describing relatively simple or uniform functionality, are too abstract to properly describe entities as complex as e-commerce services or feature rich telecommunications services. The web services community has partially acknowledged this, as description languages like WSCL and OWL-S have enriched interface information with additional fragments of component semantics. In this paper, we naturally extend this progression by proposing that services be ...


Keywords: behavioral model, executable specification, execution monitoring, service discovery, validation

4 Fast algorithm for creating space efficient dispatching tables with application to multi-dispatching



Yoav Zibin, Joseph Yossi Gil

November 2002 **ACM SIGPLAN Notices , Proceedings of the 17th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications**, Volume 37 Issue 11

Full text available:  pdf(312.23 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The dispatching problem can be solved very efficiently in the single-inheritance~(SI) setting. In this paper we show how to extend one such solution to the multiple-inheritance~(MI) setting. This generalization comes with an increase to the space requirement by a small factor of κ . This factor can be thought of as a metric of the complexity of the topology of the inheritance hierarchy. On a data set of ~35 hierarchies totaling some~64 thousand types, our dispatching data structure, based on a ...

5 Efficient topology-aware overlay network



Marcel Waldvogel, Roberto Rinaldi

January 2003 **ACM SIGCOMM Computer Communication Review**, Volume 33 Issue 1

Full text available:  pdf(583.22 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Peer-to-peer (P2P) networking has become a household word in the past few years, being marketed as a work-around for server scalability problems and "wonder drug" to achieve resilience. Current widely-used P2P networks rely on central directory servers or massive message flooding, clearly not scalable solutions. Distributed Hash Tables (DHT) are expected to eliminate flooding and central servers, but can require many long-haul message deliveries. We introduce Mithos, an content-addressable overl ...

6 A semi-automatic approach to home video editing



Andreas Girgensohn, John Boreczky, Patrick Chiu, John Doherty, Jonathan Foote, Gene Golovchinsky, Shingo Uchihashi, Lynn Wilcox

November 2000 **Proceedings of the 13th annual ACM symposium on User interface software and technology**

Full text available:  pdf(1.06 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: automatic video clip extraction, video analysis, video editing, video exploration

7 User modeling I: What would they think?: a computational model of attitudes



Hugo Liu, Pattie Maes

January 2004 **Proceedings of the 9th international conference on Intelligent user interface**

Full text available:  pdf(350.99 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A key to improving at any task is frequent feedback from people whose opinions we care about: our family, friends, mentors, and the experts. However, such input is not usually available from the right people at the time it is needed most, and attaining a deep understanding of someone else's perspective requires immense effort. This paper introduces a technological solution. We present a novel method for automatically modeling a person's attitudes and opinions, and a proactive interface called "Wh ...

Keywords: affective interfaces, affective memory, user modeling

8 [Placing search in context: the concept revisited](#)



Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, Eytan Ruppin

April 2001 **Proceedings of the tenth international conference on World Wide Web**

Full text available:  pdf(235.96 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: context, invisible web, search, semantic processing, statistical natural language processing

Results 1 - 8 of 8

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

 SEARCH

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Placing search in context: the concept revisited

Full text Pdf (236 KB)

Source [International World Wide Web Conference archive](#)
Proceedings of the tenth international conference on World Wide Web [table of contents](#)
 Hong Kong, Hong Kong
 Pages: 406 - 414
 Year of Publication: 2001
 ISBN:1-58113-348-0

Authors [Lev Finkelstein](#) Zapper Technologies Inc., 3 Azrieli Center, Tel Aviv 67023, Israel
[Evgeniy Gabrilovich](#) Zapper Technologies Inc., 3 Azrieli Center, Tel Aviv 67023, Israel
[Yossi Matias](#) Zapper Technologies Inc., 3 Azrieli Center, Tel Aviv 67023, Israel
[Ehud Rivlin](#) Zapper Technologies Inc., 3 Azrieli Center, Tel Aviv 67023, Israel
[Zach Solan](#) Zapper Technologies Inc., 3 Azrieli Center, Tel Aviv 67023, Israel
[Gadi Wolfman](#) Zapper Technologies Inc., 3 Azrieli Center, Tel Aviv 67023, Israel
[Eytan Ruppín](#) Zapper Technologies Inc., 3 Azrieli Center, Tel Aviv 67023, Israel

Sponsors [SIGWEB](#): ACM Special Interest Group on Hypertext, Hypermedia, and Web
[SIGLINK](#): Hypertext, Hypermedia, and Web
 IW3C2 : International World Wide Web Conference Committee

Publisher ACM Press New York, NY, USA

Additional Information: [references](#) [citations](#) [index terms](#) [collaborative colleagues](#) [peer to peer](#)

Tools and Actions: [Discussions](#) [Find similar Articles](#) [Review this Article](#)
[Save this Article to a Binder](#) [Display Formats: BibTex](#) [EndNote](#)

DOI Bookmark: Use this link to bookmark this Article: <http://doi.acm.org/10.1145/371920.372094>
[What is a DOI?](#)

↑ REFERENCES

Note: OCR errors may be found in this Reference List extracted from the full text article. ACM has opted to expose the complete List rather than only correct and linked references.

1 [Krishna Bharat, SearchPad: explicit capture of search context to support Web search, Proceedings of the 9th international World Wide Web conference on Computer networks : the international journal of computer and telecommunications netowrking: the international journal of computer and telecommunications netowrking, p.493-501, June 2000, Amsterdam, The Netherlands](#)

2 [Jay Budzik , Kristian J. Hammond, User interactions with everyday applications as context for just-in-time information access, Proceedings of the 5th international conference on Intelligent user interfaces, p.44-51, January 09-12, 2000, New Orleans, Louisiana, United States](#)

3 D. Butler. Souped-up search engines. Nature, Vol. 405, pp.112-115, May 2000

4 R.O. Duda and P.E. Hart. Pattern Classification and Scene Analysis. New York: John Wiley and Sons, 1973.